

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problems Mailbox.**

(12) **UK Patent Application** (19) **GB** (11) **2 323 946** (13) **A**

(43) Date of A Publication 07.10.1998

(21) Application No 9706821.7

(22) Date of Filing 04.04.1997

(71) Applicant(s)
Sony United Kingdom Limited
 (Incorporated in the United Kingdom)
 The Heights, Brooklands, WEYBRIDGE, Surrey,
 KT13 0XW, United Kingdom

(72) Inventor(s)
Martin Rex Dorricott

(74) Agent and/or Address for Service
D Young & Co
 21 New Fetter Lane, LONDON, EC4A 1DA,
 United Kingdom

(51) INT CL⁶
G06F 17/30

(52) UK CL (Edition P)
G4A AJR AUDB
U1S S2106

(56) Documents Cited
WO 93/17393 A1 WO 90/13089 A1 US 4899292 A

(58) Field of Search
 UK CL (Edition O) **G4A AJR AUDB**
 INT CL⁶ **G06F 17/30**
 On-line: **WPI**

(54) Abstract Title

Decompressing images retrieved from a database and storing them in a cache

(57) To facilitate retrieval of footage from a video archive, a database is used whose entries comprise images representing the items in the video archive. When the results of a query are returned, the user reviews th images and selects those corresponding to the desired video item(s). The database images are stored in a compressed format and when search results are provided to a user they are decompressed and cached for display.




Fig. 2




Fig. 1




Fig. 2

DATABASE ACCESSING METHOD AND APPARATUS

This invention relates to methods and apparatus for accessing databases, particularly databases having database records including image data.

5 When a database is accessed, a typical arrangement is that a "query" or search statement is passed to the database, and in response the database returns details of any "hits" - i.e. those database entries or records matching the criteria specified in that query.

10 Some databases have records containing image data. An example is a video archiving database, where single images or "picture stamps" characterising the whole of an archived video clip can be stored in a database alongside textual comments and keywords relating to that clip. When the database responds to a user query, the user can be presented with a scrollable display of the images and the textual information, enabling the user to identify a desired database record quickly.

15 However, when the database records contain image data, this response from the database can represent a large amount of data to be transferred to the client application or client computer making the query. This is because digitised images generally contain a large amount of data in comparison to textual information.

20 This large amount of data to be transferred can be acceptable when the querying application and the database are resident on the same host computer, but when they are resident on separate machines linked, for example, by a bandwidth-limited network connection such as an ethernet link, transferring such large amounts of data can become a problem. The problem is made worse in situations where the database can be accessed by multiple client computers, or in which the network is
25 already heavily loaded with other data traffic.

 One apparent way of alleviating this problem of slow transfer of the digitised images would be to compress those images using a compression technique such as the known JPEG technique. However, any delay in transferring the images to the client making the database query is then added to by the time required to decompress the
30 images before they can be displayed. In the scrolling system described above, such delays can be subjectively very disturbing to the user.

 This invention provides apparatus for accessing a database having database

records comprising data-compressed image data, the apparatus comprising:

a database server apparatus operable to respond to a database query by returning database records to the client process matching the criteria specified by the database query; and

5 a client process for issuing database queries to the database server and for displaying information indicative of the database records returned by the database server apparatus;

in which:

10 the client process has a data decompressor for decompressing data-compressed images received from the database server and a data cache for storing decompressed image data representing those images;

the client process is operable to retrieve, decompress, display and store in the data cache data-compressed a subset of the images relating to the database records returned by the database server; and

15 the client process is then operable to retrieve, decompress and store in the data cache data-compressed images relating to the database records returned by the database server but not currently required for display.

20 The invention addresses the problems described above by providing a cache store of the *decompressed* image data at the client workstation. So, in addition to images required for current display to the user, other images relating to, for example, other database "hits" under the current query or to more detail of the currently displayed hits.

25 This arrangement of locally caching decompressed images can reduce the delay between a user command and the display of a corresponding image (assuming it is cached) and so avoid or reduce the subjective delays described above.

An embodiment of the invention will now be described, by way of example only, with reference to the accompanying drawings, throughout which like parts are referred to by like references, and in which:

30 Figure 1 is a schematic diagram of a video archiving system; and
Figure 2 is a schematic diagram of a client workstation in the system of Figure 1.

Figure 1 is a schematic diagram of a video archiving system.

The system comprises an client workstation 10, a device and tape control computer 20, a database server 30 and an archive controller 35, all linked by an ethernet local area network (LAN) connection, an RS-422 router 15, a video tape recorder (VTR) 40, a video (and audio) router 50, a data converter 60 and an archive tape recorder 70.

The archive controller 35 controls and coordinates the overall operation of the device and tape controller 20, and communicates with the database server (to consult or modify the database held on that server) and with the client workstation (to receive commands).

The database server holds a database of the archived material (see below). The device and tape controller 20 responds to commands issued by the archive controller to generate specific RS-422 or SCSI commands to drive the controlled devices 40, 50, 60, 70 (and possibly other similar devices - not shown).

The RS-422 router is operable to route RS-422 commands to the VTR from either the device and tape controller 20 or the client workstation 10.

Finally, each client workstation is available for use by a human archivist in preparing material for archive and corresponding entries in the database and for browsing through and retrieving database entries or records. The client workstations are in effect the "client", in that they communicate high level archiving tasks to the archive controller, which then schedules and carries out those tasks - e.g. by instructing the device and tape controller 20.

The archiving system works in two main modes of operation. When incoming video material (with or without accompanying audio material) is to be archived, the material is replayed from the VTR 40 and viewed at the client workstation 10. A human operator (referred to as the archivist) then assigns a title, various keywords, an abstract and other information to the video material. This additional information forms a database entry which is stored in a database maintained on the database server 30. In the meantime, the video material is stored on the archive tape recorder 70. This "archive creation" process will be described in more detail below.

In a second ("archive retrieval") process, the human operator uses the client workstation 10 (or another client workstation connected to the network) to browse through the database entries held on the database server 30. For example, the user

might be searching for video material having a particular associated keyword. When the user finds one or more pieces of video material in this manner which he believes are appropriate to the current requirements, he can schedule a retrieval operation to retrieve that video material from the archive tape recorder.

5 These two processes will now be described in more detail.

Archive Creation

10 In this mode of operation, incoming video material is replayed from the VTR (for example, a Sony Digital BetaCam video tape recorder) and is encoded into a known compressed video format (the so-called MPEG II 4:2:2 P@ML format). The compressed video is supplied in a known data transmission format known as "SDDI" via the router 50 to the client workstation 10, where it can be viewed by the archivist.

15 At this archive creation stage, the routing applied by the video router 50 is controlled by the device and tape control computer 20, which in turn responds to commands by the client workstation 10. However, for rapid control by the human operator, the RS-422 commands which control the VTR 40 are routed, via the RS-422 router, directly from the client workstation 10 to the VTR 40. So, as a first step, the device and tape controller 20 uses an RS-422 control line 25 to set up the RS-422 router to patch RS-422 commands from the client workstation to the VTR, and from then on during the archive creation process the commands are passed directly. Thus, 20 if the archivist wishes to replay a particular portion of the tape, he makes the appropriate commands at the client workstation 10, which are communicated via RS-422 router to the VTR 40 using a conventional RS-422 control protocol.

25 The archivist replays through the material to be archived, perhaps several times over to gain a full understanding and familiarity with the events covered by the video material. The archivist is presented with software-generated database entry forms on the screen of the client workstation 10, at which the archivist can enter information to be used in classifying that video material. The archivist can also select individual frames from the material to act as "picture stamps", to be stored in the database record as images particularly characterising that video material or a camera shot within the video material. These items of information together form a database entry to be added to a conventional database maintained on the database server 30.

30 When the database entry has been prepared and stored, the video material is

transferred from the VTR 40, via the router 50 to the data converter 60. In doing this, the archive controller instructs the device and tape controller to set up the relevant devices, including setting up the RS-422 router now to route RS-422 commands from the device and tape controller to that VTR, and then to issue the necessary operating commands to set the VTR into play mode and the archive tape recorder into a record mode. The data converter 60 receives the video data in the SDDI synchronous transmission format and outputs the video data to the archive tape recorder 70 in a SCSI data transfer format.

Because the nature of the SCSI transmission format does not tend to allow data transfer operations to be initiated and started quickly enough for a synchronous video data transfer, the data converter 60 contains a large buffer memory (not shown) to provide any necessary delay to the data being passed to or from the data converter 60.

The archive tape recorder itself is a helical scan digital tape recorder such as the Sony GY10 tape recorder. This device is primarily a "data" recorder, and so will accept data other than video data, but it is noted that in this embodiment the bulk of the material stored on the archive tape recorder is in fact video data.

The video data and certain other data is then recorded on the archive tape recorder 70. Specifically, the video data is routed as described above from the VTR 40, via the router 50 to the data converter 60 and finally to the archive tape recorder 70; and the other associated data (e.g. data derived from the database entry for that material) is routed from the database server 30 to the device and tape controller 20 via the LAN, and then from the device controller 20 to the archive tape recorder 70 via the SCSI link.

In a typical application, there is a very large amount of video data to be archived in this way, so a robotic library system may be used in association with the archive tape recorder 70. In any event, whether a robotic or a manual library system is used, it is very likely that a large number of tapes will be required, and so the database entry for the video material archived on a particular tape (which is stored in the database server 30) also contains a tape identifier which identifies the appropriate tape (or tapes, if multiple copies are made) on which that material is stored.

Archive Retrieval

The first step in an archive retrieval operation is to identify the appropriate piece of video material to be retrieved. The user can do this by browsing through the database stored on the database server 30 from a computer workstation (such as the client workstation 10) connected via the ethernet link to the database server 30. This searching mechanism may be conventional, using key words, title information or abstract text to search for a required piece of video material. The search or query process is described in more detail below.

Once the required material has been identified, the user can issue a command which is transmitted via the ethernet link to the archive controller, which in turn issues a command to the device and tape control computer 20 to retrieve the required video data from the archive tape recorder 70. Again, executing these commands will include setting up the SCSI link between the archive tape recorder and the data converter, setting up the RS-422 links between the other devices, and issuing correctly timed commands to initiate play and record modes of the devices.

In general, in a typical system these commands will have to be queued by the archive controller and then initiated and executed in a priority order.

When a command is executed to retrieve data from the archive tape recorder, it is first necessary to have the correct tape placed in the archive tape recorder 70 (e.g. by the robotic library system - not shown). Then, the device and tape control computer 20 sets up the SCSI link from the archive tape recorder to the data converter and issues appropriate RS-422 commands to the data converter, the router and the VTRs 40 so that the video data is replayed from the archive tape recorder 70, communicated via the SCSI link to the data converter 60, converted into an SDDI synchronous transmission format, and then supplied via the router 50 to the VTR 40 for recording. The material can then be viewed, either during this process or later from the VTR, at one of the computer workstations.

Thus, in the example operation described above, it may be necessary to issue RS-422 commands simultaneously to the data converter, the router and to the VTRs. As well as the requirement for issuing the commands simultaneously, a more onerous requirement is that the commands often have to be timed at sub-field accuracy with respect to an external source of video syncs. This requirement can arise during an

archive creation process when an incoming video signal from another part of the studio is to be archived, or during an archive retrieval process when an archived section of video material has to be retrieved at a particular time for recording or transmission elsewhere in the studio.

5 The servers and controllers (i.e. the device and tape controller 20, the database server 30 and the archive controller 35) used in the apparatus of Figure 1 are standard PC-compatible general purpose computers, operating under the Microsoft Windows NT operating system (for the database server 30 and the archive controller 35) or the QNX real-time operating system available from the QSSL company (for
10 the device and tape controller 20). These operating systems allow message-driven multitasking within each server between multiple concurrent processing tasks, which provides efficient use of the server's processing capacity. However, this does mean that the time at which a particular operation takes place is not generally predictable to the accuracy of a fraction of a video frame period required for issuing control
15 commands to the controlled video devices 40, 50, 60.

Figure 2 is a schematic diagram showing the operation of the client workstation 10 when involved in a database browsing operation.

5 The client workstation 10 comprises a display and user interface controller 100, a cache controller 110, a data decompression device 120, a cache storage device 130 and a display and input device 140 (e.g. a computer keyboard and a computer display screen).

At the start of the database browsing operation, the display and user interface controller 100 sends a user query to the database server 30. This is a standard database accessing technique: the query might be a keyword for comparison with an
10 abstract associated with video material held on the database server 30, a time period (e.g. "find all material originated between 1 and 3 March 1997"), a subject reference or similar. In response, the database returns details of a number of database "hits" - i.e. database records complying with the search criteria set out in the query - together with image "handles" corresponding to those database hits.

15 The image handles are pointers to image data stored in the database representing picture stamps associated with the database entries. Digitally storing an image, even using data compression, generally requires significant amounts of data

to be stored. If all this data was returned to the display and user interface controller 100 in response to a database query, there could be a significant delay involved. This is why only the image "handles" or pointers are returned to the display and user interface controller, as the image handles occupy very much less data than a complete image.

5 The display and user interface controller selects those images which it requires immediately for display in the currently selected output format on the display device and sends an image request to the cache controller 110. The cache controller 110 checks in a cache register 115 whether the required images are held in the cache 130 and, if not, issues a query back to the database server 130 to obtain the required image data. The image data is returned to the cache controller.

10 The image data is in a compressed form, and so the cache controller passes the compressed image data it receives from the database to the decompression device 120, with the decompressed data being stored in the cache 130. The decompressed image data for the currently-required images is also returned via the cache controller to the display and user interface controller 100 for display on the display and input device 140.

Once the immediate needs of the display and user interface controller have been addressed in this way, the display and user interface controller then initiates the downloading and decompression of other images relating to the same database search. An example of this process is given below, but in general terms the operation is as above, in that the display and user interface controller 100 issues further requests to the cache controller for images to be cached (but not immediately supplied back to the display and user interface controller). The cache controller requests those images from the database server 30 and, after decompression by the decompression device 120, stores them in the cache 130. The cache controller also updates the cache register 115 with details of the newly cached images.

Example

30

In the example, the database stored on the database server 30 has database records each having the following format:

Main text description of archived video material clip	text relating to sub-section 1	text relating to sub-section 2	. . .	text relating to sub-section 4
Main picture stamp characterising entire clip	picture stamp characterising sub-section 1	picture stamp characterising sub-section 2	. . .	picture stamp characterising sub-section 4

5

In this format, the database record corresponding to a material clip includes an overall description and an overall ("main") characterising picture stamp. Then individual sub-sections of the material - perhaps different camera shots within the material - are provided with individual descriptions and picture stamps.

10

Assume that the user carries out a search which uncovers 12 records as database "hits" - i.e. there are 12 records which comply with the criteria specified in the search query.

At an initial level, the user might want to see only the main picture stamp and the overall description from each record.

15

If this information from all 12 database hits can be displayed simultaneously, then the display and user interface controller would request the cache controller to obtain and decompress the image data from all 12 main picture stamps straight away.

20

However, if all 12 records cannot fit onto the display screen at once, they can be presented in a scrollable format so that a subset of the records is visible at any time. In this case, the display and user interface controller would first request the cache controller to obtain and decompress the subset of the 12 main images for immediate display, and then would request the cache controller to obtain and decompress the remaining ones of the 12 main images to allow a smooth and rapid scroll through the set of 12 hits. For example, if information relating to four hits can be displayed at a time, the display and user interface controller would initiate retrieval, decompression and caching of the main picture stamps for the first four records (for immediate display), and would then initiate retrieval, decompression and caching of the remaining eight main picture stamps.

25

30

Once the "main" images relating to the database hits have been retrieved and decompressed in this way, the display and user interface controller instructs the cache

controller to obtain and decompress the images relating to sub-sections of each database hit. This operation is in anticipation of a request by the user to look in more detail at one of the database records which came up as a hit in the search. So, if and when the user does make such a request, the intention of this arrangement is that the required picture stamp images for the sub-sections of that record will by then have been retrieved from the database, decompressed and stored in the cache 130.

At all stages, the cache controller 110 maintains the cache register indicating which images are cached in the cache 130, together with a time at which the images were cached.

The cache is capable of holding a particular number of decompressed images - for example, in this embodiment 100 images can be cached. When a new image is to be cached in a cache 130 which is already full, then the oldest cached image is discarded unless that image relates to a record currently under examination by the user.

CLAIMS

1. Apparatus for accessing a database having database records comprising data-compressed image data, the apparatus comprising:
- 5 a database server apparatus operable to respond to a database query by returning database records to the client process matching the criteria specified by the database query; and
- a client process for issuing database queries to the database server and for displaying information indicative of the database records returned by the database server apparatus;
- 10 in which:
- the client process has a data decompressor for decompressing data-compressed images received from the database server and a data cache for storing decompressed image data representing those images;
- 15 the client process is operable to retrieve, decompress, display and store in the data cache data-compressed a subset of the images relating to the database records returned by the database server; and
- the client process is then operable to retrieve, decompress and store in the data cache data-compressed images relating to the database records returned by the database server but not currently required for display.
- 20
2. Apparatus according to claim 1, in which the client process runs on a client data processing apparatus different to the database server apparatus, the client data processing apparatus being connected to transfer data to and from the database server apparatus by a data communication link.
- 25
3. Apparatus according to claim 1 or claim 2, in which the database records comprise a main image characterising the record and one or more sub-images characterising sub-sections of the record, the client process being operable to retrieve, decompress and store in the data cache the data-compressed main images relating to the database records returned by the database server but not currently required for display, and then to retrieve, decompress and store in the data cache the data-
- 30

compressed sub-images relating to the database records returned by the database server.

5 4. Apparatus according to any one of the preceding claims, in which, when the data cache has been filled with image data, the client process is operable to overwrite the earliest-stored image in the data cache to store a newly retrieved and decompressed image, unless that earliest-stored image relates to a database record currently being displayed.

10 5. Apparatus for accessing a database, the apparatus being substantially as hereinbefore described with reference to the accompanying drawings.



Application No: GB 9706821.7
Claims searched: All

Examiner: Matthew Gillard
Date of search: 3 June 1997

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:
UK CI (Ed.O): G4A AUDB, AJR
Int CI (Ed.6): G06F 17/30
Other: On-line: WPI

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
X	WO 93/17393 A1 (SALIENT). See page 21.	1 at least
X	WO 90/13089 A1 (ARCHIVAGE SYSTEMES). See page 5 & ll. 14-19 on page 16.	1 at least
X	US 4899292 (IMAGE STORAGE). See col. 7, ll. 48-62; col. 8, l. 47 - col. 9, l. 17; col 10, l. 67 - col. 11, l. 12.	1 at least

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.